

RIACS FINAL REPORT

October 1, 1996 through September 30, 1997

Contract NAS 2-96027

Submitted to:

Contracting Officer: Catharine P. Levin
Code IC, MS 241-1

Submitted by:

Research Institute for Advanced Computer Science (RIACS)

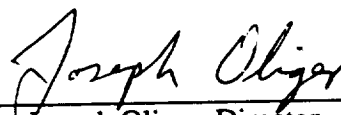
An Institute of:

Universities Space Research Association (USRA)

RIACS Principal Investigator
Joseph Oliger

NASA Technical Monitor
Mr. Anthony Gross

COTR
Mr. Anthony Gross


Joseph Oliger, Director

100-6118
100-6118
22-1
094560

TABLE OF CONTENTS

I. INTRODUCTION	2
II. RESEARCH PROJECTS	4
A. High Performance Computing	4
B. High Performance Networks	14
III. TECHNICAL REPORTS	17
IV. PUBLICATIONS	21
V. SEMINARS AND COLLOQUIA	24
VI. OTHER ACTIVITIES	26
VII. RIACS STAFF	27

L INTRODUCTION

Joseph Oliger, Director

The Research Institute for Advanced Computer Science (RIACS) was established by the Universities Space Research Association (USRA) at the NASA Ames Research Center (ARC) on June 6, 1983. RIACS is privately operated by USRA, a consortium of universities with research programs in the aerospace sciences, under contract with NASA. RIACS performs computer science research in collaboration with NASA scientists to solve challenging scientific problems in support of NASA's goals and missions. RIACS serves as an intermediary between the NASA Ames Research Center and the academic community. Research is carried out by a staff of full-time scientists, augmented by visitors, students, post doctoral candidates and visiting university faculty.

The Ames Research Center has recently been designated NASA's Center of Excellence in Information Technology. In this capacity, Ames has been charged with the responsibility to build an information technology research program that is preeminent within the Agency. Accordingly, RIACS has recently reorganized its activities.

The primary mission of RIACS is chartered to carry out research and development in computer science. This work is devoted in the main to tasks that are strategically enabling with respect to NASA's bold missions in space exploration and aeronautics. There are three foci for this work:

- *high-performance computing*
- *cognitive and perceptual prostheses (computational aids designed to leverage human abilities)*
- *autonomous systems*

An objective of RIACS is to broaden the base of researchers working in these areas of importance to the nation's aeronautics and space enterprises. In this connection, RIACS works to foster collaborative links between scientists at Ames and RIACS' staff and visitors. In particular, through its visiting scientist program, RIACS facilitates the participation of university-based researchers, from the U.S and abroad in this research and development.

In 1997, RIACS had 3 staff scientists, 8 visiting scientists, 1 post doctoral scientist, 8 consultants, 3 research associates and 1 system administrator.

During this report period Professor Wei-Pai Tang of the University of Waterloo, Professor Marsha Berger of New York University, Professor Tony Chan of UCLA, Associate Professor David Zingg of University of Toronto, Professor Robert MacCormack of Stanford, Professor Eli Turkel of Tel Aviv University, Professor James Sethian of University of Cal Berkeley and Assistant Professor Andrew Sohn of New Jersey Institute of Technology have been visiting RIACS.

RIACS held two seminars during this report period. The seminars were held to discuss needs and opportunities in basic research in computer science in and for NASA applications.

*Topic: Unstructured Grid Applications and General Preconditioning in CFD
Date: July 22, 1997*

There were 6 talks given by university scientists. Part 1 discussions of the seminar gave an overview of the various subtopics in Unstructured grid applications; (i) Progress in parallel Schur complement preconditioning for computational fluid dynamics (ii) Agglomerative multilevel methods for elliptic problems on unstructured grids and (iii) Elliptic multilevel solvers on unstructured grids. Part 2 discussions consisted of subtopics on Preconditioning Techniques: (i) Advances in preconditioning techniques, (ii) Local and global preconditioning in computational algorithms for aerodynamic flows and (iii) High level one-way dissection preconditioners for unsteady incompressible Navier-Stokes flow.

*Topic Level Set Methods: Evolving Interfaces in Geometry, Fluid Mechanics, Computer Vision and Materials Sciences.
Date: April 10 - 21, 1997*

The speaker, James Sethian, Professor, Univ of Cal Berkeley, presented a tutorial on level set methods, which are mathematical and numerical techniques for tracking propagating interfaces. These techniques naturally handle sophisticated interface motion, including the generation of corners and cusps, topological changes of merging and breaking, and complex evolutions in three dimensions and higher. They have been successfully applied to a wide range of problems, including two-fluid interfaces and mixing, combustion, image processing and computer vision, medical imaging, grid generation, computation of first arrival times in seismic events and in robotic path planning, shape recognition, and etching and deposition simulations in the fabrication of microelectronic components. The tutorial will cover all of these applications, as well as the details of the numerical methodology and implementation.

In addition to RIACS seminars, RIACS also participated in the NASA Ames Open House for the public in September 1997.

RIACS technical reports are usually preprints of manuscripts that have been submitted to research journals or conference proceedings. A list of these reports for the period January 1, 1997 through September 30, 1997 is in the Reports and Abstracts section of this report.

II. RESEARCH PROJECTS

A. HIGH PERFORMANCE COMPUTING

PARALLEL LOAD BALANCER FOR ADAPTIVE UNSTRUCTURED MESHES

Leonid Oliker, Rupak Biswas and Roger C. Strawn (US Army AFDD)

Dynamic mesh adaption on unstructured grids is a powerful tool for computing solutions of unsteady three-dimensional problems that require grid modifications to efficiently resolve solution features. An efficient parallel implementation of these methods is extremely difficult to achieve, primarily due to the load imbalance created by the dynamically changing non-uniform grid. To address this problem, we have developed PLUM, an automatic and portable framework for performing load-balanced adaptive large-scale numerical computations in a message-passing environment.

During FY97, we completed the implementation and integration of all major components within our dynamic load-balancing strategy. This includes interfacing a parallel solution-adaptive procedure to a fast repartitioner and an efficient data remapper. Previous results indicated that mesh repartitioning and data remapping are potential bottlenecks for performing large-scale flow computations. We resolve these issues and demonstrate that our framework scales with the number of processors.

Our load-balancing procedure has five novel features. (i) A dual graph representation of the initial computational mesh that keeps the complexity and connectivity constant during the course of an adaptive computation. (ii) The integration of a parallel mesh repartitioning algorithm avoids a potential serial bottleneck. Five state-of-the-art schemes from the MeTis and JOSTLE packages were examined. Results indicate that for certain classes of unsteady adaption, globally repartitioning the computational mesh produces higher quality results than diffusive repartitioning schemes. (iii) Optimal and heuristic remapping algorithms quickly assign partitions to processors so that the redistribution cost is minimized. (iv) An efficient data movement scheme allows remapping and mesh subdivision at a significantly lower cost than previously reported. (v) An accurate cost metric predicts the data remapping time, by considering both the interprocessor communication overhead and the computational cost of data reshuffling. This cost measure is then compared to the computational gain that would be achieved with a balanced workload to determine the viability of the data redistribution, and hence the load-balancing step.

The mesh adaption and global load-balancing codes are written in C, C++, and MPI. The effectiveness of our strategy has been verified on both a steady state helicopter rotor problem and an unsteady shock wave simulation. Finally, we are examining portability by comparing results on the three vastly different architectures of the IBM SP2, the SGI/Cray T3E, and the SGI/Cray Origin2000.

DEVELOPMENT OF A 3D UNSTRUCTURED GRID CODE BASED ON A FINITE VOLUME FORMULATION AND APPLIED TO THE NAVIER-STOKES EQUATIONS

Michel Delanaye

The aim of the present research carried out during the report period, 1997, was the development of a high-order accurate finite volume scheme for solving the 3D Navier-Stokes equations on unstructured hybrid grids.

Current state of the art techniques for simulating flows with unstructured grids are essentially based on the calculations of second-order accurate advective fluxes through faces of the control volumes. A fundamental property that should be fulfilled by a scheme intended to be used to simulate high-Reynolds number flows is that the truncation error of the advective term discretization should not exhibit any second order derivatives. Indeed, in this case, the leading truncation error term produces some spurious artificial dissipation that can spoil the physical diffusive effects. This fundamental property should also hold whatever grid distortions, which is crucial for unstructured grids that intend to be always distorted. On the other hand, the design of high-order accurate techniques is very important to decrease the mesh size and constraints on mesh quality.

In our recent Ph.D. thesis, we showed that a third-order accurate calculation of the advective fluxes is required to achieve this fundamental property. The scheme we propose is based on a third-order accurate reconstruction of the flow variables in each control volume, and on a third-order accurate integration of the numerical fluxes (Roe Riemann solver) along the edges or faces of the control volume in 2D or 3D respectively. The application of such a high-order accurate scheme in 3D is very challenging because of the demanding cost, but also because of the related robustness problems.

Unstructured grids are most of the time considered as made of simplices: triangles in 2D and tetrahedra in 3D. Indeed, those elements allow to cover the domain nearly automatically. However, regarding the simulation of high-Reynolds number flows, the use of simplices is not really adequate for resolving the very different length scales present in the flow. The discretization of a boundary layer with tetrahedra yields a very large number of distorted elements, which is not suitable for the cost of the calculation, accuracy and sometimes robustness. For that purpose, we have considered unstructured grids as a collection of four different types of elements: tetrahedra, prisms, pyramids and cubes. Such grids are often referred to as hybrid. It provides us with the maximum flexibility in the grid generation process, allowing to obtain "good" grids with "good" connectivities and a "good" distribution of nodes in crucial regions like boundary layers, while maintaining the possibility of easily generating grids for complex configurations. Finite volume schemes based on a second-order calculation of the fluxes are based on a dual control volume approach. In this case, the degrees of freedom are associated to the vertices of the mesh. For unstructured grids made of tetrahedra, this choice is optimum because the alternative of storing the degrees of freedom at the centroid of the control volume (cell-centered approach) would require about 5 times more the memory. The cost of the dual control volume approach is essentially proportional to number of edges of the mesh. Indeed, a composite dual face of a dual control volume is associated to each edge of the mesh. That composite dual face is actually a collection of small triangular faces. In the case of a second-order accurate calculation of the flux, the composite dual face is considered as a single face and an approximation of the flux can be calculated by a simple edge based formula. In order to obtain a third-order approximation of the flux, we now have to consider the flux through each small triangular faces composing the dual face (median dual control volume). It can be easily shown that the cost of the third-order method is now proportional to 12 times the number of cells, which is much larger than the number of edges.

For third-order method, the dual control volume approach is therefore not suitable, and a more classical cell-centered approach is preferred despite the increased number of degrees of freedom and associated memory.

The main part of the design of a third-order accurate method is actually the reconstruction of the data in each control volume. In our method, this is achieved by a truncated third-order Taylor series expansion around the centroid of the control volume. That expansion involves the calculation of first and second order derivatives. The procedure consists in using information from the surrounding cells to calculate those derivatives with a prescribed accuracy. A two-step least square method is employed. In the first step, the first derivatives are calculated to first-order, in the second step, the second derivatives are calculated to first-order and the previously calculated first derivatives are corrected to reach the second-order accuracy. This procedure has been demonstrated to show better accuracy. It is also flexible because different stencils can be used for the first and second derivatives. The choice of the stencil is crucial to achieve accuracy and robustness. A sufficient number of "well located" nodes has to be devised. The use of too many nodes is detrimental to the cost, and too few nodes will not yield the required accuracy. Unfortunately, we have not been able to find a "universal" rule for choosing the stencil for arbitrary 3D meshes. However, we have found some good combinations for the four different type of elements.

In order to quickly obtain steady state solutions, we use a pseudo unsteady approach based on a fully implicit scheme. At each pseudo time-step, a linearization of the system is performed (Newton method). The resulting Jacobian is known to be indefinite and non-symmetrical. We therefore use the robust GMRES algorithm to solve this linear system. It is used in its finite difference version which avoids the actual calculation and storage of the Jacobian matrix of the high-order scheme. The GMRES algorithm is preconditioned by an ILU(0) decomposition of an approximate Jacobian based on a first-order scheme. The routines used in the implicit scheme are based on the PETSc library developed at the Argonne National Lab.

Preliminary results have been obtained for inviscid flow calculations. The third-order finite volume scheme has been tested for the simulation of the inviscid flow in a channel with one wall perturbed by a sine bump. The results show the improved accuracy with respect to a more classical second-order method. Subsonic flows around wings have also been computed by the code using very general hybrid grids.

CARTESIAN GRID METHODS FOR COMPLEX GEOMETRY

Marsha Berger, Michael Aftosmis(NASA Ames) and John Melton (NASA Ames)

In this approach, a solid object is superimposed on an underlying Cartesian grid, and the flow is computed around the object. This makes the problem of volume grid generation substantially easier, with the bulk of the work reduced to finding intersections between a possibly complex configuration and a regular Cartesian grid. However, the difficulty of grid generation is traded for the difficulties in the flow solver of imposing solid wall boundary conditions on a non-body fitted grid. Our previous work on flow solvers for this kind of grid however indicates that acceptable results that maintain second order accuracy over the entire flow field can be obtained.

Our research this year has focused on the robustness and efficiency of Cartesian mesh generation. We have developed algorithms (borrowing greatly from the computational geometry literature), in an effort to make the grid generator as robust as possible. Since the purpose of this approach is to automate the grid generation process and handle the very complex cases, we have carefully designed the various steps to avoid the usual pitfalls associated with other grid generation methods (along with the usual ways to handle them, such as jiggling the

mesh, inflate the geometry, etc.) Especially for use in a time dependent setting (e.g. one with moving geometry), these other techniques can not be applied.

Our mesh generation process begins with watertight triangulations of each component in a configuration. This approach helps alleviate the burden on the CAD operator, since the triangulations need not be constrained to the intersection curves between components, and neighboring components are not required to have commensurate length scales. The triangulations are pre-processed to extract the wetted (exposed) surface of the configuration. This removes the possibility of internal geometry, i.e. geometry that is in fact internal to another component, which greatly slowed the mesh generation procedures. We have also developed a uniform way of handling the degenerate cases (for example where two triangles intersect on one of their edges). These cases typically take 90% of the coding though they occur 1% of the time. For example, component definitions typically end at the symmetry plane, where a lot of degeneracy's can be found. For these cases we are using the "Simulation of Simplicity" approach by [Edelsbrunner and Mucke, ACM Trans. Graphics, 9(1), Jan. 1990], which breaks the degeneracy's using a virtual displacement of the data.

The main idea is to perturb the data into a non-degenerate position based on perturbations of a particular form. An asymptotic expansion of the new data around the old value of the determinant will perturb the exact zero to one sign or the other.

MULTIGRID METHODS FOR SOLVING ELLIPTIC PROBLEMS ON UNSTRUCTURED GRIDS

Susie Go, Tony Chan (UCLA) and Timothy Barth (NASA Ames).

Unstructured grids are easily adapted to complex geometry's and steep gradients in the solution, thus their increasing popularity. They are, unfortunately, not naturally suited for multilevel methods since these methods require a grid hierarchy upon which to define coarse grid problems--something which is not available when using unstructured grids.

We have been working on developing robust domain decomposition and multigrid methods for solving elliptic problems on unstructured grids. In particular, we have looked at ways to properly define the subspace problems for node-nested multilevel methods when the physical boundaries of coarse grids do not match the boundary of the fine grid problem. We have shown that with proper treatment of boundary conditions, the multilevel methods can achieve optimal convergence rates on elliptic problems. The current work was done using a library of basic linear and non-linear solvers which was developed at Argonne National Labs, known as the Portable Extensible Toolkit for Scientific Computing (PETSc). PETSc was chosen because it is a currently supported library which has both sequential- and parallel-processor (using MPI message passing) capability.

Extension of these multilevel methods to more complicated equations such as the Euler equations, is being done to see if similar effects occur. This phase of the project required the development of software for fluid flow. The Element Library for Fluid Flow (ELF) was written with Timothy Barth. The components currently available in the ELF library are finite element discretization of a system of equations by Galerkin, Galerkin Least Squares and Discontinuous Galerkin methods with a choice for piecewise constant, linear, and quadratic functions, as well as several different quadrature rules. The library is being integrated with the PETSc solvers to provide a host of different solvers for two- or three-dimensional Euler flow.

Using the same definition for coarse grid problems, another application of the multilevel solvers is to use them to solve eigenvalue problems. One particular use is the eigenvalue problem

which arises in spectral partitioning methods. We use the full approximation schedule multigrid method developed for eigenproblems by Brandt, McCormick, Ruge (1983). The new features which make this implementation different from other multilevel partitioners is in the definition of the intergrid operators for defining the coarse problem as well as the possibility for convergence proofs.

ALGEBRAIC NON-OVERLAPPING DOMAIN DECOMPOSITION METHODS FOR COMPRESSIBLE FLUID FLOW PROBLEMS ON UNSTRUCTURED MESHES

Tony Chan, Tim Barth (NASA Ames) and Wei-Pei Tang (U. Waterloo)

We consider preconditioning methods for convection dominated fluid flow problems based on a non-overlapping Schur complement domain decomposition procedure for arbitrary triangulated domains. The triangulation is first partitioned into a number of subdomains and interfaces which induce a natural 2X2 partitioning of the p.d.e. discretization matrix. We view the Schur complement induced by this partitioning as an algebraically derived coarse space approximation. This avoids the known difficulties associated with the direct formation of an effective coarse discretization for advection dominated equations. By considering various approximations of the block factorization of the 2X2 system, we have developed a family of robust preconditioning techniques. These approximations are introduced to improve both the sequential and parallel efficiency of the method without significantly degrading the quality of the preconditioner. The specific approximations that we have used include ILU-preconditioned GMRES subdomain solves, localized approximation of the interface Schur complement, and limited level-fill ILU interface backsolves. A computer code based on these ideas has been developed and tested on the IBM SP2 using MPI message passing protocol. A number of 2-D CFD calculations will be presented for both scalar advection-diffusion equations and the Euler equations. These results show very good scalability of the preconditioner as the number of processors is increased while the number of degrees of freedom per processor is fixed.

NUMERICAL METHODS FOR THE COMPRESSIBLE NAVIER-STOKES EQUATIONS WITH APPLICATIONS TO AERODYNAMIC FLOWS

David Zingg

David Zingg continued his collaborative work with Dr. Tom Pulliam on Computational algorithms for the Navier-Stokes equations applied to aerodynamic flows. Topics studied included Newton-Krylov methods, the convective upstream split pressure scheme (CUSP), and boundary schemes for higher-order methods. Uniformly second-order boundary schemes (leading to third-order global accuracy) are now implemented and working in two flow solvers for the Navier-Stokes equations, one incompressible, the other compressible. In addition, studies of the convective upstream split pressure (CUSP) scheme have continued, with a new soft limiter showing good performance. Overall CUSP leads to comparable accuracy to matrix dissipation at a reduced cost. An extended abstract based on this work has been submitted to the 29th AIAA Fluid Dynamics Conference. Dr. Zingg gave a presentation entitled "Progress in Computational Algorithms for Aerodynamic Flows," at the NASA Ames Research Center.

RESEARCH IN AERODYNAMIC SHAPE OPTIMIZATION

James Reuther

Since the inception of CFD, researchers have sought not only accurate aerodynamic prediction methods for given configurations, but also design methods capable of creating new optimum configurations. Yet, while flow analysis can now be carried out over quite complex configurations using the Navier-Stokes equations with a high degree of confidence, direct CFD based design is still a daunting challenge for complex three-dimensional problems. This is especially true in problems where viscous effects play a dominant role. The main effort of this research is the introduction of new technology to overcome the difficulties present in traditional aerodynamic optimization methods. The CFD-based aerodynamic design methods of the past can be grouped into two basic categories: inverse methods, and numerical optimization methods.

Inverse methods derive their name from the fact that they invert the goal of the flow analysis algorithm. Instead of obtaining the surface distribution of an aerodynamic quantity, such as pressure, for a given shape, they calculate the shape for a given surface distribution of an aerodynamic quantity. Most of these methods are based on potential flow techniques, and few of them have been extended to three-dimensions. The common trait of all inverse methods is their computational efficiency. Typically, transonic inverse methods require the equivalent of 2-10 complete flow solutions in order to render a complete design. Since obtaining a few solutions for simple two-dimensional and three-dimensional designs can be done in at most a few hours on modern computer systems, the computational cost of most inverse methods is considered to be minimal. Unfortunately, they suffer from many limitations and difficulties the most glaring of which is that the objective is built directly into the design process and thus cannot be changed to an arbitrary or more appropriate objective function.

A traditional alternative, which avoids some of the difficulties of inverse methods while incurring a heavy computational expense, is the use of numerical optimization methods. The essence of these methods is straightforward: a numerical optimization procedure is coupled directly to an existing CFD analysis algorithm. The numerical optimization procedure attempts to extremize a chosen aerodynamic measure of merit which is evaluated by the chosen CFD code. Most of these optimization procedures require gradient information in addition to evaluations of the objective function. Here, the gradient refers to changes in the objective function with respect to changes in the design variables. The simplest method of obtaining gradient information is by finite differences. In this technique, the gradient components are estimated by independently perturbing each design variable with a finite step, calculating the corresponding value of the objective function using CFD analysis, and forming the ratio of the differences. These methods are very versatile, allowing any reasonable aerodynamic quantity to be used as the objective function. They can be used to mimic an inverse method by minimizing the difference between target and actual pressure distributions, or may instead be used to maximize other aerodynamic quantities of merit such as L/D . Unfortunately, these finite difference numerical optimization methods, unlike the inverse methods, are computationally expensive because of the large number of flow solutions needed to determine the gradient information for a useful number of design variables. Tens of thousands of flow analyses would be required for a complete three-dimensional design.

In this research, a new method is developed that avoids the limitations and difficulties of traditional inverse methods while retaining their inherent computational efficiency. The method dramatically reduces the cost of aerodynamic optimization by replacing the expensive finite-difference method of calculating the required gradients with an adjoint variable formulation. After deriving the differential form of the adjoint equations and posing the correct boundary

conditions based on the objective function, the resulting system is discretized and solved on the same mesh as that used for the flow solution. A significant economization is thus achieved by applying the same subroutines used for the flow solution to the solution of the adjoint equations. The resulting design process requires only one flow calculation and one adjoint calculation per gradient evaluation, as opposed to the hundreds required for a finite-difference gradient involving hundreds of design variables. In practice the computational cost of the new method is two orders of magnitude less than a conventional approach. Considerable effort has been focused in the last two years to develop control theory-based aerodynamic shape optimization methods. This effort has been conducted by a team of researchers from around the nation whose major contributors include Prof. Antony Jameson of Stanford University, Prof. Luigi Martinelli of Princeton University, Prof. Juan J. Alonso of Stanford University, Dr James Farmer and myself. Many of the core subroutines upon which the research has been formulated is the intellectual property of Intelligent Aerodynamics International. The work that has taken place in the last three years can be broken down into three specific areas.

- A) Two-dimensional and three-dimensional proof-of-concept studies.
- B) The development and demonstration of a three-dimensional research tool for complex configurations.
- C) NASA and industrial evaluation and feedback.

During the first year, work was primarily focused in area (1) and to a lesser extent areas (2) and (3). At the beginning of this program at RIACS, methods were in place which showed that control theory could be used in conjunction with numerical optimization and computational fluid dynamics to create efficient design tools for flows governed by the potential flow equation (AIAA Paper 94-0499).

During the course of the first year of the program the development of adjoint methods was extended to treat the Euler equations. In our paper at the Multi-Disciplinary Optimization conference during summer 1994 (AIAA paper 94-4272, also RIACS report 94.18), results were shown demonstrating that control theory based on the Euler equations could be used to design airfoils that operate under transonic conditions. Various objective functions were demonstrated showing the versatility of the new method. In the work presented at VKI, the first examples of three-dimensional wing design using control theory were presented. Finally, in a paper presented at the January 1995 Aerospace Sciences Meeting (AIAA paper 95-0123, also RIACS report 95.01) results for the design of wing and wing-body configurations over general meshes were shown.

One of the dramatic successes in the first year involved the participation of Beechcraft Aircraft Division of Raytheon, Inc. Raytheon entered into a cooperative agreement with NASA Ames Research Center to explore the usefulness of the adjoint-based design optimization methods. Between March and May of 1995, a team of scientists from Raytheon and Ames were able to combine their talents by employing a preliminary version of the three-dimensional design code, described in RIACS report 95.01, to develop a new business jet wing for the Premier I configuration. This one-month design of a new transonic wing contrasts with the usual development time of more than a year for traditional methods. Raytheon has since conducted wind tunnel tests confirming that the new wing design realizes its predicted performance and launched the design for production. They subsequently took 51 orders for the new airplane on the day the design was announced. Furthermore, Raytheon has been so impressed by the capability of adjoint-based design methods that they are now incorporating them into their own aircraft design environments. A paper authored by both NASA and Raytheon personnel that presents the basic design strategy and its outcome was presented at the Aerospace Sciences Meeting, January 1996 (AIAA paper 96-0554, also RIACS 96.03).

Another group that has taken a keen interest in our research from its first year is the NASA High Speed Research Program (HSR) group. In their effort to create economically viable supersonic transport configurations for the next century they are investigating the use of aerodynamic shape optimization to improve aerodynamic performance. Both the traditional as well as adjoint-based design methods studied by our group at Ames have been tested by the HSR community. A paper that gives an example of the capabilities of this emerging technology for supersonic design was presented at the American Society of Mechanical Engineers annual winter meeting in November 1995 (also RIACS 95.14).

At the beginning of the second year of this research, experience with both Raytheon and the HSR group highlighted the need to develop an enhanced implementation of the aerodynamic shape optimization method which would allow the treatment of more complex geometry's. Until that time, only a single-grid-block design method had been developed, capable of handling wing/body configurations but leaving engine nacelle effects to be modeled by approximations at best. However, the real-life problems presented by industry required the design method to handle complete aircraft configurations, which in turn mandated either an extension to a multiblock grid topology or a switch to unstructured meshes. The former path was chosen because of its relatively straightforward implementation and the natural avenue which multiple blocks provide towards a parallelization of the process.

The first paper demonstrating the new multiblock capability was presented at the 34th Aerospace Sciences Meeting (AIAA paper 96-0094, RIACS report 96.02). Following this paper, the focus quickly turned to a parallel form of the multiblock software. This was essential because the added complexity of complete aircraft configurations required a significant increase in the harnessed computer power. A paper presented at the Multi-disciplinary Optimization Conference in September 1996 highlighted this parallel multiblock capability.

The second year was also characterized by a significant effort to enhance the software for the HSR work. Since details of this work cannot be presented in view of its sensitive nature, it must suffice to state that both the single-block and multiblock codes were modified so that HSR-specific design problems could be treated robustly and efficiently. One major activity was the incorporation of a constrained optimization capability as opposed to the use of an unconstrained algorithm. This was necessitated by the hundreds of geometric constraints imposed on the HSRP configurations (such as on wing spar thickness', fuel volume, and cabin dimensions). The year culminated in the successful application of the HSRP-specific versions of the software to an industry-established test-bed configuration. Independent methods were also applied to the same problem by Boeing and McDonnell Douglas teams. This constrained optimization exercise showed the software used at NASA Ames to be effective and favorable. Further efforts during this last year have focused on enhancing the multiblock design capability to treat even more realistic problems. The first step in this path was the inclusion of constraints and the treatment of multiple design points. A paper presented at the 35th Aerospace Sciences Meeting (AIAA paper 97-0103, RIACS report 97.02) highlighted these capabilities. The next step was the inclusion of a viscous design capability through the extension of the underlying flow solver from the Euler equations to the Navier-Stokes equations. An important point is that the solution cost in terms of computer time to solve the Navier-Stokes equations as opposed to the Euler equations is a factor of about 5. Further, since multiple flow solutions are required to solve a design problem, the use of parallel computing, first introduced in the second year of this research, has become an essential capability of the software. Navier-Stokes-based design problems typically require on the order of 32 SGI Origin 2000 CPUs for roughly 24 hours. This level of computer resources corresponds to computer turn-around times that would be unacceptable on the fastest serial CPUs that are available today. Considerable time was thus invested to ensure that the software ran robustly in parallel and on various platforms. To date, the parallel multiblock code has been ported to the IBM SP2, the CRAY J90 and C90, and SGI Origin 2000, and a cluster of HP workstations. The need for parallel efficiency on these

widely-varying architectures required careful management and tuning of the interprocessor communication costs. Issues included load balancing, bandwidth minimization, latency reduction, and scalability with respect to total mesh size. A paper presented at the AIAA 13th Biannual CFD Conference (AIAA paper 97-1893 RIACS report 97.05) discussed both the extension to the Navier-Stokes equations and the details related to improving the parallel performance.

Since this most recent paper, attention has once again reverted to practical applications of the software. We are currently involved in testing the enhanced multiblock method on HSRP problems as well as on cooperative projects with both the newly-merged Boeing/McDonnell Douglas company and the Raytheon/Beechcraft company. One necessary modification to the design method that took place this year was the inclusion of engine inflow and outflow boundary conditions such that propulsion induced effects could be accounted for during the course of the design. This capability is set to be exercised in the upcoming year. Further research is also continuing on several fronts to advance the technology of aerodynamic shape optimization. One area of attention is the treatment of viscous design problems. To date the parallel multiblock code only has the algebraic Baldwin-Lomax turbulence model fully implemented thus work is proceeding to develop a structure within the code which permits the choice from amongst an entire suite of turbulence models. The first two turbulence models currently under study are the Spalart-Allmaras model and the K-Omega SST model of Menter. In parallel with these developments in turbulence model implementations, work is underway to include an integral boundary layer method into the design code as an alternative to switching to the Navier-Stokes equations. This would have an advantage in reducing the computational expense of a design run but will only be appropriate for certain attached flow design problems. So far work in this area has been limited to two dimensions but will expand to treat three dimensions in the coming year.

Even with much work still to be accomplished it is nevertheless gratifying that the developments achieved thus far have demonstrated beyond a doubt the great value of adjoint-based aerodynamic design. It is hoped that with all of these advances, the greater aeronautical science community will in the future adopt these new ideas into their production design environments. Certainly if the work in conjunction with Raytheon is any indication, this is already taking place.

S-HARP: A PARALLEL DYNAMIC SPECTRAL PARTITIONER

Andrew Sohn

Computational science problems based on adaptive meshes involve dynamic load balancing when implemented on parallel machines. This dynamic load balancing requires frequent partitioning of computational meshes at runtime. This report presents a parallel dynamic partitioner, called S-HARP. The underlying principles of S-HARP are the fast feature of inertial partitioning and the quality feature of disconnectivity-based partitioning. S-HARP partitions a graph from scratch, requiring no partition information from previous iterations. Two types of parallelism have been exploited in S-HARP, fine-grain loop-level parallelism and coarse-grain recursive parallelism. The parallel partitioner has been implemented in Message Passing Interface on Cray T3E and IBM SP2. Experimental results indicate that S-HARP can partition the mesh of over 100,000 vertices into 256 partitions in 0.2 seconds on a 64-processor T3E. S-HARP is much more scaleable than ParaMeTiS1.0, giving over 15-fold speedup on 64 processors while ParaMeTiS1.0 gives a few-fold speedup. Experimental results demonstrate that S-HARP is three to 10 times faster than the dynamic partitioners ParaMeTiS and Jostle.

NUMERICAL SCHEMES FOR THE HAMILTON-JACOBI AND LEVEL SET EQUATIONS ON TRIANGULATED DOMAINS

James Sethian and Tim Barth (NASA Ames)

Borrowing from techniques developed for conservation law equations, we developed numerical schemes which discretize the Hamilton-Jacobi (H-J), level set, and Eikonal equations on triangulated domains. The first scheme developed is a provably monotone discretization for certain forms of the H-J equations. Unfortunately, the basic scheme lacks proper Lipschitz continuity of the numerical Hamiltonian. By employing a "virtual" edge flipping technique, Lipschitz continuity of the numerical flux is restored on acute triangulations. Next, schemes were developed based on the weaker concept of positive coefficient approximations for homogeneous Hamiltonians. These schemes possess a discrete maximum principle on arbitrary triangulations and naturally exhibit proper Lipschitz continuity of the numerical Hamiltonian. Finally, a class of Petrov-Galerkin approximations were invented. These schemes are stabilized via a least-squares bilinear form. The Petrov-Galerkin schemes do not possess a discrete maximum principle but generalize to high order accuracy. Discretization of the level set equation also requires the numerical approximation of a mean curvature term. A simple lumped-Galerkin approximation was then developed and analyzed using maximum principle analysis. The use of unstructured meshes permits several forms of mesh adaptation which have been incorporated into numerical examples. These numerical examples include discretizations of convex and nonconvex forms of the H-J equation, the Eikonal equation, and the level set equation.

The impact of this work is as follows: this research develops a general methodology for treating level set methods and the more general Hamilton-Jacobi equations on triangulated domains. This opens up the possibility of adaptive mesh refinement techniques for propagating interfaces, including boundary-fitted internal boundary conditions and high resolution. Future work should include applying this work to a host of interface problems, including those in semi-conductor manufacturing, and materials sciences.

APPLICATION OF HIGH-ORDER SHOCK CAPTURING SCHEMES TO DIRECT SIMULATION OF TURBULENCE

Neil Sanham and Dr. Helen C. Yee (NASA Ames)

The purpose of this visit was to continue an investigation into the applicability of high-order shock-capturing schemes to direct numerical simulation of turbulence. On an earlier visit (January 1997) several methods had been programmed into a Navier-Stokes code and two test cases had been developed, one of vortex pairing in a Mach 0.8 mixing layer and the other of an oblique shock wave impacting on a free shear layer. This work has been continued on the latest visit by running the program for the two test cases, with additional debugging tests and code optimisations.

The methods use fourth order Runge-Kutta time advancement with compact and non-compact schemes of up to sixth order. Differentiation routines were validated with test functions and the convergence of different methods to the same solution on fine grids was verified. The compact and non-compact TVD schemes were optimised for the C90 computer. For the vortex pairing case a total of 12 simulations were run to compare the various schemes. A preliminary conclusion is that a second and fourth-order dissipation extension of the shock capturing schemes is required in order to achieve benefits from the higher order schemes. Two simulations of the shock-wave/shear-layer interaction test case on a fine grid were run, also showing an improved solution with a fourth order non-compact method compared to an earlier second-order method. Overall computational cost is also reduced due to a less restrictive stability criterion for the time-step.

B. HIGH PERFORMANCE NETWORKS

MULTICAST TECHNOLOGY

Marjory J. Johnson

Multicast is the transmission of data from a single source to multiple receivers. IP Multicast is a critical networking technology for NASA, enabling applications in several areas, e.g., distributed computing, distributed simulation, collaborative design and analysis, and video conferencing.

M. Johnson spent considerable time this past year learning all aspects of multicast technology, because of its fundamental importance for NASA networking projects. She is conducting a survey of various approaches for achieving reliable multicast, and is developing a framework for evaluating them within the context of NASA applications. A paper is in progress.

IP multicast traffic is based on the UDP networking protocol, whereas the majority of today's Internet traffic is based on TCP. While effective flow control and congestion control techniques have been developed and refined for TCP traffic, UDP traffic is not subject to these TCP window-control mechanisms. Because of the expected large volume of multicast traffic in future networks, this is a potentially serious problem. Another paper in progress evaluates various approaches for controlling congestion for multicast applications.

Current plans include experimental activities with the NASA Research and Education Network to test analytical results regarding congestion control and Quality of Service (QoS) provision for multicast applications.

NETWORK TESTBEDS

Marjory J. Johnson

M. Johnson has actively participated in the planning of major networking testbeds, including the Next Generation Internet, the NASA Research and Education Network, and the National Transparent Optical Network. These testbeds will serve as vehicles for development of new networking technologies and services that are required to enable future applications, as well as for demonstration of these new applications.

NEXT GENERATION INTERNET (NGI)

The Next Generation Internet (NGI) is a new 3-year, \$300 million federal initiative that will "create the foundation for networks of the 21st century." The specific objectives of the initiative are: 1) to create a network infrastructure connecting selected universities and national laboratories that is 100 to 1000 times faster than the current Internet, 2) to develop the technology to enable demanding new networking applications that support important national goals and missions, e.g., scientific research, national security, distance education, environmental monitoring, and health care, and 3) to demonstrate these new applications. The Workshop on Research Directions for the Next Generation Internet, an invitation-only two-day workshop, was convened in May 1997 to plan the research agenda needed to accomplish these goals. M. Johnson was invited to participate in the workshop upon acceptance of her white paper entitled "Some Quality of Service Issues." All of the accepted white papers, which provided a catalyst for discussion at the workshop and which are now part of the formal record of the workshop, are available at <http://www.cra.org/Policy/NGI/accpapers.html>.

At the workshop M. Johnson was assigned to a working group on quality of service. Other working groups included architecture, applications, middleware, security, and traffic engineering. Documents prepared by the various working groups, completed after the workshop adjourned, are included in the workshop report, "Research Challenges for the Next Generation Internet," published by the Computing Research Association and available electronically at http://www.cra.org/main/research_chall.pdf.

NASA Ames is the lead center for the agency's portion of the NGI project.

NASA RESEARCH AND EDUCATION NETWORK (NREN)

The NASA Research and Education Network (NREN) program forms the core of the NASA NGI program.

M. Johnson participated in the Second Annual High Performance Computing and Communications/NASA Research and Education Network (HPCC/NREN) Workshop. M. Johnson was a member of the Advanced Aerospace Design working group. The other working groups included Astrobiology, Astrophysics, Earth Sciences, Telemedicine, and Space Exploration. The primary objective of the workshop was to identify applications for each of the above disciplines. In the Advanced Aerospace Design working group we identified design environments, virtual facilities, and physics-based deep analysis as three major application areas, and listed several future applications within each area. Then we identified enabling networking technologies and specific network capabilities to support the above applications. Development of collaborative work environments and tele-socialization received prominent notice. Results from the workshop will be posted on the NREN web site, <http://www.nren.nasa.gov>.

M. Johnson is collaborating with NREN personnel at Ames to develop mechanisms for providing Quality of Service to enable real-time applications on NREN.

NATIONAL TRANSPARENT OPTICAL NETWORK (NTON)

The National Transparent Optical Network is a 10-gigabit-per-second fiber-optic ring around the Bay Area, currently connecting Lawrence Livermore National Lab, UC Berkeley, Pacific Bell - San Ramon, and Sprint - Burlingame. Connection to NASA Ames is underway. The stated objectives of the NTON Consortium are "to create an open, all-optical network that demonstrates critical wave-division-multiplexing technologies and control strategies required for terabit per second optical networks for US DoD, Information Industry and consumer communications."

M. Johnson participated in several meetings to plan NASA Ames's role in the National Transparent Optical Network (NTON) testbed and in an extended testbed that is formed by hosting a ground terminal for the NASA ACTS satellite at one of the NTON sites. The NTON/ACTS extended testbed will reach NASA centers on the east coast. We have identified seven or eight potential applications, representing key activities at Ames. We are also planning research projects using NTON.

One application that we are pursuing for the testbed is to run some aircraft design code (provided by James Reuther) on a workstation cluster distributed via NTON between NASA Ames and Sandia - Livermore.

SUPERCOMPUTER CONSOLIDATION PROJECT

Marjory J. Johnson

M. Johnson is working with NREN personnel and researchers at the University of California - San Francisco (UCSF) to select a suitable simulation package to support the NASA supercomputer consolidation project in evaluating alternatives for locating the agency's supercomputer facilities. We are evaluating several simulation packages in the context of modeling the supercomputer workload from all the NASA centers. The CACI COMNET III simulation package has been installed at UCSF. M. Johnson is experimenting with using this package, but access delay between Ames and the workstation at UCSF hosting the simulation package makes the task difficult. A complementary simulation package will be selected for installation at Ames.

BAY AREA GIGABIT NETWORK (BAGNet) DATA ANALYSIS

Marjory J. Johnson

Before the Bay Area Gigabit Network (BAGNet) was disbanded a year ago, Bellcore captured data on the network using a tool that they had developed for use with ATM networks. RIACS participated in some of these data-capture sessions, and M. Johnson is analyzing selected subsets of that data. The first data set was collected while several sites were generating multicast streams; the second was collected while running an image browsing application that was developed at RIACS.

M. Johnson is investigating patterns of packet loss for the browser application, particularly when cells from multiple applications are interleaved. Results from this analysis provide insight into the problem of providing adequate congestion control for future applications.

MISCELLANEOUS ACTIVITIES

Marjory J. Johnson

M. Johnson has been active in the general networking community, by serving on various program committees and by refereeing papers for journals. She served on a review panel for a DoE laboratory technology program in November of 1996. She is also a member of the ISO/TC20/USSCAG 13 committee to develop communication standards for space missions.

III. TECHNICAL REPORTS

97.01 HARP: A Fast Spectral Partitioner

Horst D. Simon (NERSC - LBL), Andrew Sohn (NJIT) and Rupak Biswas (MRJ Technology Solutions) March 1997 (10 pages)

Appeared in the 9th ACM Symposium on Parallel Algorithms and Architectures, Newport, Rhode Island, June 1997

Partitioning unstructured graphs is central to the parallel solution of computational science and engineering problems. Spectral partitioners, such recursive spectral bisection (RSB), have proven effective in generating high-quality partitions of realistically-sized meshes. The major problem which hindered their wide-spread use was their long execution times. This paper presents a new inertial spectral partitioner, called HARP. The main objective of the proposed approach is to quickly partition the meshes at runtime in a manner that works efficiently for real applications in the context of distributed-memory machines. The underlying principle of HARP is to find the eigenvectors of the unpartitioned vertices and then project them onto the eigenvectors of the original mesh. Results for various meshes ranging in size from 1000 to 100,000 vertices indicate that HARP can indeed partition meshes rapidly at runtime. Experimental results show that our largest mesh can be partitioned sequentially in only a few seconds on an SP2 which is several times faster than other spectral partitioners while maintaining the solution quality of the proven RSB method. A parallel MPI version of HARP has also been implemented on IBM SP2 and Cray T3E. Parallel HARP, running on 64 processors SP2 and T3E, can partition a mesh containing more than 100,000 vertices into 64 subgrids in about half a second. These results indicate that graph partitioning can now be truly embedded in dynamically-changing real-world applications.

97.02 Constrained Multipoint Aerodynamic Shape Optimization Using an Adjoint Formulation and Parallel Computers

James Reuther, Antony Jameson (Princeton University), J Alonso (Princeton University), M. Rimlinger (Sterling Software) and D. Saunders (Sterling Software)
January 1997 (26 pages)

Presented at the AIAA 35th Aerospace Sciences Meeting and Exhibit, AIAA paper 97-0103

An aerodynamic shape optimization method that treats the design of complex aircraft configurations subject to high fidelity computational fluid dynamics (CFD), geometric constraints and multiple design points is described. The design process will be greatly accelerated through the use of both control theory and distributed memory computer architectures. Control theory is employed to derive the adjoint differential equations whose solution allows for the evaluation of design gradient information at a fraction of the computational cost required by previous design methods. The resulting problem is implemented on parallel distributed memory architectures using a domain decomposition approach, an optimized communication schedule, and the MPI (Message Passing Interface) standard for portability and efficiency. The final result achieves very rapid aerodynamic design based on a higher order CFD method.

In order to facilitate the integration of these high fidelity CFD approaches into future multi-disciplinary optimization (MDO) applications, new methods must be developed which are capable of simultaneously addressing complex geometries, multiple objective functions, and geometric design constraints. In our earlier studies we coupled the adjoint based design formulations with unconstrained optimization algorithms and showed that the approach was effective for the aerodynamic design of airfoils, wings, wing-bodies, and complex aircraft configurations. In many of the results presented in these earlier works, geometric constraints

were satisfied either by a projection into feasible space or by posing the design space parameterization such that it automatically satisfied constraints. Furthermore, with the exception of reference where the second author initially explored the use of multipoint design in conjunction with adjoint formulations, our earlier works have focused on single point design efforts. Here we demonstrate that the same methodology may be extended to treat complete configuration designs subject to multiple design points and geometric constraints. Examples are presented for both transonic and supersonic configurations ranging from wing alone designs to complex configuration designs involving wing, fuselage, nacelles and pylons.

97.03 Efficient Load Balancing and Data Remapping for Adaptive Grid Calculations

Leonid Oliker and Rupak Biswas (MRJ Technology Solutions)

April 1997 (10 pages)

Appeared in the 9th ACM Symposium on Parallel Algorithms and Architectures, Newport, Rhode Island, June 1997

Mesh adaption is a powerful tool for efficient unstructured-grid computations but causes load imbalance among processors on a parallel machine. We present a novel method to dynamically balance the processor workloads with a global view. This paper presents, for the first time, the implementation and integration of all major components within our dynamic load balancing strategy for adaptive grid calculations. Mesh adaption, repartitioning, processor assignment, and remapping are critical components of the framework that must be accomplished rapidly and efficiently so as not to cause a significant overhead to the numerical simulation. Previous results indicated that mesh repartitioning and data remapping are potential bottlenecks for performing large-scale scientific calculations. We resolve these issues and demonstrate that our framework remains viable on a large number of processors.

97.04 CFD Analysis and Design Optimization Using Parallel Computers

L. Martinelli (Princeton University), J.J. Alonso (Princeton University), A. Jameson (Princeton University) and James Reuther
January 1997 (38 pages)

A versatile and efficient multi-block method is presented for the simulation of both steady and unsteady flow, as well as aerodynamic design optimization of complete aircraft configurations. The compressible Euler and Reynolds Averaged Navier-Stokes (RANS) equations are discretized using a high resolution scheme on body-fitted structured meshes. An efficient multigrid implicit scheme is implemented for time-accurate flow calculations. Optimum aerodynamic shape design is achieved at very low cost using an adjoint formulation. The method is implemented on parallel computing systems using the MPI message passing interface standard to ensure portability. The results demonstrate that, by combining highly efficient algorithms with parallel computing, it is possible to perform detailed steady and unsteady analysis as well as automatic design for complex configurations using the present generation of parallel computers.

97.05 An Efficient Multiblock Method for Aerodynamic Analysis and Design on Distributed Memory Systems

James Reuther, J. J. Alonso (Princeton University), J. C. Vassberg (Douglas Aircraft Co.), Antony Jameson (Princeton University) and L. Martinelli (Princeton University)
January 1997 (27 pages)

The work presented in this paper describes the application of a multiblock gridding strategy to the solution of aerodynamic design optimization problems involving complex configurations. The design process is implemented in parallel using the MPI (Message Passing Interface) Standard such that it can be efficiently used on a variety of distributed memory systems ranging from traditional parallel computers to networks of workstations. Substantial improvements to the parallel performance of the baseline method are developed, with particular attention to their impact on the scalability of the program as a function of the mesh size. Drag minimization calculations at a fixed coefficient of lift are presented for a business jet configuration that includes wing, pylon, aft-mounted nacelle, and vertical and horizontal tails. An aerodynamic design optimization is performed with both the Euler and Reynolds Averaged Navier-Stokes (RANS) equations governing the flow solution and the results are compared. These sample calculations establish the feasibility of efficient aerodynamic optimization of complete aircraft configurations using the RANS equations as the flow model. There still exists, however, the need for detailed studies of the importance of a true viscous adjoint method which holds the promise of tackling the minimization of not only the wave and induced components of drag, but also the viscous drag.

97.06 Dynamics of Numerics and Spurious Behaviors in CFD Computations

Helen C. Yee (NASA Ames Research Center) and Peter K. Sweby (University of Reading)
June 1997 (148 pages)
An invited review paper for Journal of Computational Physics

The global nonlinear behavior of finite discretizations for constant time steps and fixed or adaptive grid spacings is studied using tools from dynamical systems theory. Detailed analysis of commonly used temporal and spatial discretizations for simple model problems is illustrated. The role of dynamics in the understanding of long time behavior of numerical integration and the nonlinear stability, convergence, and reliability of using time-marching approaches for obtaining steady-state numerical solutions in computational fluid dynamics (CFD) is exploited. The study is complemented with examples of spurious behavior observed in steady and unsteady CFD computations. The CFD examples were chosen to illustrate non-apparent spurious behavior that was difficult to detect without extensive grid and temporal refinement studies and some knowledge from dynamical systems theory. Studies revealed the various possible dangers of misinterpreting numerical simulation of realistic complex flows that are constrained by available computing power. In large scale computations where the physics of the problem under study is not well understood and numerical simulations are the only viable means of solution, extreme care must be taken in both computation and interpretation of the numerical data. The goal of this paper is to explore the important role that dynamical systems theory can play in the understanding of the global nonlinear behavior of numerical algorithms and to aid the identification of the sources of numerical uncertainties in CFD.

97.07 Runge-Kutta Methods for Linear Differential Equations

David W. Zingg and Todd T. Chisholm (University of Toronto Institute for Aerospace Studies)
 July 1997 (15 pages)

Three new Runge-Kutta methods are presented for numerical integration of systems of linear inhomogeneous ordinary differential equations (ODEs) with constant coefficients. Such ODEs arise in the numerical solution of the partial differential equations governing linear wave phenomena. The restriction to linear ODEs with constant coefficients reduces the number of conditions which the coefficients of the Runge-Kutta method must satisfy. This freedom is used to develop methods which are more efficient than conventional Runge-Kutta methods. A fourth-order method is presented which uses only two memory locations per dependent variable, while the classical fourth-order Runge-Kutta method uses three. This method is an excellent choice for simulations of linear wave phenomena if memory is a primary concern. In addition, fifth- and sixth-order methods are presented which require five and six stages, respectively, one fewer than their conventional counterparts, and are therefore more efficient. These methods are an excellent option for use with high-order spatial discretizations.

97.08 Load Balancing Sequences of Unstructured Adaptive Grids

Rupak Biswas (MRJ Technology Solutions) and Leonid Oliker
 July 1997, (6 pages)

Mesh adaption is a powerful tool for efficient unstructured grid computations but causes load imbalance on multiprocessor systems. To address this problem, we have developed (\sf PLUM), an automatic portable framework for performing adaptive large-scale numerical computations in a message-passing environment. This paper makes several important additions to our previous work. First, a new remapping cost model is presented and empirically validated on an SP2. Next, our load balancing strategy is applied to sequences of dynamically adapted unstructured grids. Results indicate that our framework is effective on many processors for both steady and unsteady problems with several levels of adaption. Additionally, we demonstrate that a coarse starting mesh produces high quality load balancing, at a fraction of the cost required for a fine initial mesh. Finally, we show that the data remapping overhead can be significantly reduced by applying our heuristic processor reassignment algorithm.

97.09 Repartitioning and Load Balancing Adaptive Meshes

Rupak Biswas (MRJ Technology Solutions) and Leonid Oliker
 September 19, 1997, (23 pages)
 presented at IMA Workshop on Grid Generation and Adaptive Algorithms, 5/97 , Minneapolis, MN

Mesh adaption is a powerful tool for efficient unstructured-grid computations but causes load imbalance on multiprocessor systems. To address this problem, we have developed PLUM, an automatic portable framework for performing adaptive large-scale numerical computations in a message-passing environment. This paper presents several experimental results that verify the effectiveness of PLUM on sequences of dynamically adapted unstructured grids. We examine portability by comparing results between the distributed-memory system of the IBM SP2, and the Scalable Shared-memory MultiProcessing (S2MP) architecture of the SGI/Cray Origin2000. Additionally, we evaluate the performance of five state-of-the-art partitioning algorithms that can be used within PLUM. Results indicate that for certain classes of unsteady adaption, globally repartitioning the computational mesh produces higher quality results than diffusive repartitioning schemes. We also demonstrate that a coarse starting mesh produces high quality load balancing, at a fraction of the cost required for a fine initial mesh. Finally, we show that the data redistribution overhead can be significantly reduced by applying our heuristic processor reassignment algorithm to the default partition-to-processor mapping given by partitioners.

IV. PUBLICATIONS

Michel Delanaye, Geuzaine, Ph and Essers, J.A., "Development and Application of Quadratic Reconstruction Schemes for Compressible Flows on Unstructured Adaptive Grids", AIAA paper 97-2120, 13th AIAA CFD Conference, Snowmass, June 1997.

Michel Delanaye, Geuzaine, Ph and Liu, Y, "Application of a High-Order Reconstruction Scheme to Turbulent Flow Calculations Using Hybrid Cartesian Adaptive Unstructured Grids", AIAA paper 98-0546, to be presented at the 36th AIAA Aerospace Sciences Meeting and Exhibit, Reno, January 1998.

Michel Delanaye, "Development and Application of High-Order Accurate Finite Volume Schemes For Simulating 2D and 3D Compressible Flows on Unstructured Adaptive Grids", seminar at the Von Karman Institute, Brussels, September 1997.

Roger C. Strawn, Leonid Oliker and Rupak Biswas, "New Computational Methods for the Prediction and Analysis of Helicopter Noise," Journal of Aircraft, July 1997.

Rupak Biswas, Leonid Oliker and Andrew Sohn, "Global Load Balancing with Parallel Mesh Adaption on Distributed-Memory Systems," Proceedings of Supercomputing '96, Pittsburgh, Pennsylvania, Nov. 17-22, 1996.
<http://www.supercomp.org/sc96/proceedings/SC96PROC/BISWAS/INDEX.HTM>

Rupak Biswas and Leonid Oliker, "Load Balancing Unstructured Adaptive Grids for CFD Problems," CD-ROM Proceedings of the 8th SIAM Conference on Parallel Processing for Scientific Computing, Minneapolis, Minnesota, Mar. 14-17, 1997.

Rupak Biswas and Leonid Oliker, "Load Balancing Unstructured Adaptive Grid Computations," Abstracts of the 4th U.S. National Congress on Computational Mechanics, San Francisco, California, Aug. 6-8, 1997

Leonid Oliker and Rupak Biswas, "Dynamic Domain Decomposition for Large-Scale Adaptive Calculations," Abstracts of the 10th International Conference on Domain Decomposition Methods, Boulder, Colorado, Aug. 10-14, 1997.
[presented by Leonid Oliker]

Rupak Biswas and Leonid Oliker, "Load Balancing Sequences of Unstructured Adaptive Grids," 4th International Conference on High Performance Computing, Bangalore, India, Dec. 18-21, 1997.

Rupak Biswas and Leonid Oliker, "Experiments with Repartitioning and Load Balancing Adaptive Meshes," Proceedings of the IMA Workshop on Grid Generation and Adaptive Algorithms, Institute for Mathematics and its Applications, University of Minnesota, Minneapolis, April 28 - May 2, 1997.

Tony F. Chan, Susie Go and Ludmil Zikatanov, "Lecture Notes on Multilevel Methods for Elliptic Problems on Unstructured Grids," prepared for the lecture course "28th Computational Fluid Dynamics," Mar 3-7, 1997, von Karman Inst. for Fluid Dynamics, Belgium.
[Also appeared as UCLA Dept of Math CAM report 97-11.]

Tony F. Chan, Susie Go and Ludmil Zikatanov, "Multilevel Elliptic Solvers on Unstructured Grids," 1997 Computational Fluid Dynamics, ed. M. Hafez. Also as UCLA Dept of Math CAM report, 97-36. August 1997.

Marsha Berger, M.J. Aftosmis and J.E. Melton, "Robust and Efficient Cartesian Mesh Generation for Component-Based Geometry," AIAA Paper 97-0196, Reno, NV., Jan. 1997.

Wei-Pai Tang, "Wavelet sparse approximate inverse preconditioners", BIT 37:3 (1997), 001-017. (with T. Chan and W.L. Wan)

Wei-Pai Tang, "A fast solver for incompressible Navier-Stokes equations with finite difference methods," (with G. Golub, L. C. Huang and H. Simon), in SIAM Journal on Scientific Computing}, Accepted.

Wei-Pai Tang, J.A. George and Y. Wu, Conference Proceedings: CFD 97, May, 1997.
"Multi-level one-way dissection for unsteady incompressible Navier-Stokes flows."

James Reuther, Antony Jameson, Juan J. Alonso, Mark J. Rimlinger, and David Saunders, "Constrained Multipoint Aerodynamic Shape Optimization Using an Adjoint Formulation and Parallel Computers," AIAA (97-0103) 35th Aerospace Sciences Meeting and Exhibit, January 1997.

Luigi Martinelli, Juan J. Alonso, Antony Jameson, and James Reuther, "CFD Analysis and Design Optimization Using Parallel Computers" RIACS Technical Report 97.04 January 1997.

James Reuther, and Mark J. Rimlinger, "Development and Validation of a Multiblock Adjoint Based Design Method," HSR Configuration Aerodynamics Workshop, February 1997.

James Reuther, David Saunders, and Raymond Hicks, "Improvements to the Single-Block Adjoint Based Aerodynamic Shape Design Method SYN87-SB," HSR Configuration Aerodynamics Workshop, February 1997.

James Reuther, and Mark J. Rimlinger, "Future Advances in Aerodynamic Shape Optimization," HSR Configuration Aerodynamics Workshop, February 1997.

James Reuther, Juan J. Alonso, John C. Vassberg, Antony Jameson and Luigi Martinelli, "An efficient Multiblock Method for Aerodynamic Analysis and Design on Distributed Memory Systems," AIAA (97-1893) 13th Computational Fluid Dynamics Conference, June 1997.

PAPERS SUBMITTED TO REFEREED JOURNALS

JAMES SETHIAN

Numerical Schemes for the Hamilton-Jacobi and Level Set Equations on Triangulated Domains with T. Barth - submitted for publication, Jour. Comp. Phys.

LEONID OLIKER

Leonid Oliker and Rupak Biswas, "Performance Analysis and Portability of the PLUM Load Balancing System," Fourth World Congress on Computational Mechanics, Buenos Aires (Argentina), 29 June - 2 July 1998, submitted.

Leonid Oliker and Rupak Biswas, "Parallel Tetrahedral Mesh Adaption on the SP2," Journal of Parallel and Distributed Computing, submitted.

Leonid Oliker and Rupak Biswas, "PLUM: Parallel Load Balancing for Adaptive Unstructured Meshes," Journal of Parallel and Distributed Computing, submitted.

WEI-PAI TANG

Wei-Pai Tang, "Effective sparse approximate inverse preconditioners," submitted to SIAM Journal on Matrix Analysis and Applications.

Wei-Pai Tang, J.A. George and Y. Wu, "Multi-level one-way dissection factorization," submitted to SIAM Journal on Matrix Analysis and Applications.

V. SEMINARS AND COLLOQUIA

TONY CHAN

10th Domain Decomposition Conference, Boulder, Colorado, August 1997.
RIACS Workshop: July 24, 1997.
IMA Workshop on Parallel PDEs: June 1997.
The Von Karman Inst. lecture series, March 1997.

LEONID OLIKER

Leonid Oliker and Rupak Biswas, "Efficient Load Balancing and Data Remapping for Adaptive Grid Calculations," Proceedings of the 9th ACM Symposium on Parallel Algorithms and Architectures (SPAA), Newport, Rhode Island, June 22-25, 1997
[presented by Leonid Oliker]

Leonid Oliker and Rupak Biswas, "Dynamic Domain Decomposition for Large-Scale Adaptive Calculations," Abstracts of the 10th International Conference on Domain Decomposition Methods, Boulder, Colorado, Aug. 10-14, 1997
[presented by Leonid Oliker]

NASA Open House, NASA Ames Research Center, Moffett Field, California, September 1997.

JAMES L. REUTHER

James Reuther, and Mark J. Rimlinger, "Development and Validation of a Multiblock Adjoint Based Design Method," HSR Configuration Aerodynamics Workshop, February 1997.

James Reuther, David Saunders, and Raymond Hicks, "Improvements to the Single-Block Adjoint Based Aerodynamic Shape Design Method SYN87-SB," HSR Configuration Aerodynamics Workshop, February 1997.

James Reuther, and Mark J. Rimlinger, "Future Advances in Aerodynamic Shape Optimization," HSR Configuration Aerodynamics Workshop, February 1997.

HSR configuration aerodynamics semi-annual review/HSR Technology Configuration Aircraft design review workshop. NASA Ames Research Center, Moffett Field, California, January 1997.

HSR Configuration Aircraft workshop NASA Langley Research Center, Langley Virginia, February, 1997.

Parallel CFD 97, Machester England, May 1997.

SIAM Mini-symposium on Optimization Governed by Flow Equations at the 1997 SIAM Annual Meeting. Stanford University, Palo Alto, California, July 1997.

NASA Open House, NASA Ames Research Center, Moffett Field, California, September 1997.

Computational Aerodynamics - past present and future, A conference honoring Paul Rubbert on the occasion of his sixtieth birthday, Seattle, Washington, September 1997.

WEI- PAI TANG

Univ. of Waterloo, Institute of Computer Research Colloquium, Feb. 1997.

The Von Karman Inst. lecture series, March 1997.

IMA Workshop on Parallel PDEs: June 1997.

RIACS Workshop: July 24, 1997.

10th Domain Decomposition Conference, Boulder, Colorado, August 1997.

Second SIAM Conference on Sparse Matrices, Coeur d' Alene, Idaho, Oct, 1996.

The fifth annual Conference of the Computational Fluid Dynamics Society of Canada Victoria, British Columbia, Canada. May, 1997.

JAMES SETHIAN

Short Course in Level Set Methods, NASA Ames, March, 1997.

VI. OTHER ACTIVITIES

LEONID OLIKER

Completion of doctoral dissertation. On November 25, 1996, he successfully defended his thesis proposal at the University of Colorado. My thesis defense date is set for November 1997. My research has mainly been involved with the development of the PLUM system: a Parallel Load Balancer for Adaptive Unstructured Meshes. This framework includes the following major components: parallel mesh adaption, parallel repartitioning, optimal and heuristic processor reassignment algorithms, data remapping module, and communication modeling.

SUSIE GO

Attended SIAM Annual Meeting, July 14-18, 1997, Stanford.
Attended RIACS Workshop, July 24, 1997, NASA Ames Research Center.

TONY CHAN

I worked with Wei-Pei Tang and my Ph.D. student W.L. Wan on "Wavelet sparse approximate inverse preconditioners", BIT 37:3 (1997), 001-017.

I also worked on other unstructured grid solvers with my student Susie Go, who is being supported by RIACS, on: "Boundary treatments for multilevel methods on unstructured meshes" UCLA Dept of Math CAM report 96-30, to appear in SIAM J. Sci. Comp.

MARJORY L. JOHNSON

Participated in Workshop on Research Directions for the Next Generation Internet, Vienna, VA, May 1997.

Participated in the Second Annual High Performance Computing and Communications/NASA Research and Education Network (HPCC/NREN) Workshop, September 1997.

Member, U. S. Subcommittee Advisory Group for ISO Technical Committee 20, Subcommittee 13 (ISO/TC 20/SC 13) on Space Data and Information Transfer Systems standards.

Member, review panel for DoE laboratory technology program, November 1996.

Session chair, High Performance Networking '97 Conference.

Member of program committees for the following communications conferences: HPN '97, ECMAST '97, COST 237.

Reviewed papers for IEEE/ACM Transactions on Networks journal, International Journal of Computers and their Applications.

VII. RIACS STAFF

ADMINISTRATIVE STAFF

Joseph Oliger, Director - Ph.D., Computer Science, University of Uppsala, Sweden, 1973.
Numerical Methods for Partial Differential Equations (03/25/91 - present).

Consuelo Garza, Administrative Assistant (4/16/96 - 11/12/96).

Deanna M. Gearhart, Office Manager (2/1/96 - 8/31/97).
Administrative Assistant II (5/9/88 - 1/31/96).

Diana Martinez, Administrator (8/18/97 - present).

Steven Suhr, Systems Administrator (11/1/96 - present).

SCIENTIFIC STAFF

Dave Gehrt, JD Law, University of Washington, 1972, UNIX system administration, security,
and network based tools (1/84 - 7/85, 2/1/88 - present).

Marjory J. Johnson, Ph.D., Mathematics, University of Iowa, 1970, High-performance
networking for both space and ground applications (1/9/84 - present).

Peter J. Cheeseman, Ph.D., 1979, Artificial Intelligence, computational complexity, bayesian
inference, computer vision, plasma physics (9/1/97-present).

VISITING SCIENTISTS

Marsha Berger, Ph.D. - New York University, Computational fluid dynamics; parallel
computing (6/24/97-8/30/97).

Tony F. Chan, Ph.D. - Professor of Mathematics, University of California, Los Angeles,
Efficient algorithms in large-scale scientific computing, parallel algorithms and computational
fluid dynamics (5/30/97 and 7/21/97-7/25/97).

Andrew Sohn, Ph.D. - Assistant Professor, New Jersey Institute of Technology, Dynamic load
balancing for grid partitioning on SP-2 (6/1/97-8/31/97).

Wei-Pai Tang, Ph.D. - Professor, University of Waterloo, Canada, Numerical solution of
partial differential equations, numerical linear algebra, parallel computations (7/10/97-8/20/97).

David Zingg, Ph.D. - Associate Professor, University of Toronto, Canada, Development and
analysis of high-accuracy numerical methods applicable to simulations of fluid flows, acoustic
waves and electromagnetic waves (6/16/97-8/29/97).

Michel Delanaye, Ph.D. - Researcher in Advanced Computational Fluid Dynamics Mechanical
Engineering (4/1/97 - 9/30/97).

Robert MacCormack, Ph.D. - Professor, Stanford University, Stanford, Computational fluid dynamics, implicit numerical methods (7/29/97 - 9/30/97).

Mohamed Hafez, Lecturer - Study of convergence acceleration techniques for flow simulations (7/16/97-9/30/97).

POST-DOCTORAL SCIENTISTS

James Reuther, Ph.D. - University of California, Davis, numerical optimization aerodynamic shape optimization numerical analysis CFD(4/30/96 - present).

RESEARCH ASSOCIATES

Susie Go, MA - Applied Math, University of California, Los Angeles, multilevel methods on unstructured grids (1/21/96 - present).

Leonid Oliker, MS- Computer Science, University of Colorado, compilation of data parallel programs (9/1/94 - present).

Steven Suhr, - Computer Science, Stanford University, programming languages (7/1/92 - 10/31/96).

CONSULTANTS

Marsha Berger, Ph.D. - New York University, Computational fluid dynamics; parallel computing (1/1/93 - present).

Tony F. Chan - Professor of Mathematics, University of California, Los Angeles, Efficient algorithms in large-scale scientific computing, parallel algorithms and computational fluid dynamics (10/01/86 - present).

Richard G. Johnson, Ph.D. - Physics, Indiana University, 1956, Global environmental problems and issues (11/1/92 - present).

Neil Sanham, Ph.D. - Lecturer, Queen Mary & Westfield College, England, Direct numerical simulation of transitional and turbulent fluid flow, turbulence modeling (1/25/97 - 1/25/97 & 7/21/97-8/1/97).

Robert Schnabel, Ph.D. - Professor, University of Colorado, Boulder, Numerical computation especially optimization nonlinear equations, parallel computation (1/1/94 - present).

Jinchao Xu, Ph.D - Professor, Pennsylvania State University, numerical methods for partial differential equations, multigrid methods, parallel computations (7/21/97-7/25/97).

Wei-Pai Tang, Ph.D. - Professor, University of Waterloo, Canada, Numerical solution of partial differential equations, numerical linear algebra, parallel computations (7/1/94 -present).

Eli-Turkel, Ph.D. - Professor, Tel Aviv University, Algorithms for Navier-Stokes equations especial preconditioning for low speed flow high order accurate schemes with applications acoustics and CEM (7/21/97-7/25/97).

James Sethian, Ph.D. - Professor, University of California, Berkeley, Computational fluid mechanics, image processing, robotics and material sciences (3/1/97 - present).

Ronald Henderson, Ph.D. - Sr. Research Fellow, California Institute of Technology, Computational Fluid Dynamics, parallel computing, hydrodynamic stability, turbulence (7/20/97-8/2/97).

Dimitrios Maroudas, - Asst. Professor, Chemical Engineering, U.C. Santa Barbara, Theoretical and Computational Materials Science with emphasis on surface science and microstructure evolution in semiconductors, metallic thin films and structural alloys (9/22/97-9/23/97).